Privacy-protecting image classification within the web browser using deep learning models from Zenodo

Florian Auer, Simone Mayer, Frank Kramer

Angaben zur Veröffentlichung / Publication details:

Auer, Florian, Simone Mayer, and Frank Kramer. 2025. "Privacy-protecting image classification within the web browser using deep learning models from Zenodo." In *Intelligent health systems –from technology to data and knowledge: proceedings of MIE 2025*, edited by Elisavet Andrikopoulou, Parisis Gallos, Theodoros N. Arvanitis, Rosalynn Austin, Arriel Benis, Ronald Cornet, Panagiotis Chatzistergos, et al., 133–37. Amsterdam: IOS Press. https://doi.org/10.3233/shti250288.

Intelligent Health Systems – From Technology to Data and Knowledge
E. Andrikopoulou et al. (Eds.)

doi:10.3233/SHTI250288

Privacy-Protecting Image Classification Within the Web Browser Using Deep Learning Models from Zenodo

Florian AUER[a,1], Simone Mayer[a], and Frank KRAMER[a]
[a] *IT-Infrastructure for Translational Medical Research, University of Augsburg, Germany*
ORCiD ID: Florian Auer https://orcid.org/0000-0002-5320-8900, Simone Mayer https://orcid.org/0000-0002-7825-5738, Frank Kramer https://orcid.org/0000-0002-2857-7122

Abstract. Integrating deep learning into clinical workflows for medical image analysis holds promise for improving diagnostic accuracy. However, strict data privacy regulations and the sensitivity of clinical IT infrastructure limit the deployment of cloud-based solutions. This paper introduces WebIPred, a web-based application that loads deep learning models directly within the client's web browser, protecting patient privacy while maintaining compatibility with clinical IT environments. WebIPred supports the application of pre-trained models published on Zenodo and other repositories, allowing clinicians to apply these models to real patient data without the need for extensive technical knowledge. This paper outlines WebIPred's model integration system, prediction workflow, and privacy features. Our results show that WebIPred offers a privacy-protecting and flexible application for image classification, only relying on client-side processing. WebIPred combines its strong commitment to data privacy and security with a user-friendly interface that makes it easy for clinicians to integrate AI into their workflows.

Keywords. Deep learning, Healthcare Web Applications, Data Privacy

1. Introduction

The application of deep learning to medical imaging has shown promising results across multiple domains, including disease detection and classification. One significant barrier to widespread adoption in clinical environments is the privacy concerns surrounding sensitive patient data. Cloud-based AI solutions that send patient images to external servers for inference are often incompatible with stringent laws like General Data Protection Regulation (GDPR), which demand strict control over patient data.

To address these privacy concerns, we introduce *WebIPred*, a platform that performs image classification within the user's web browser. By leveraging pre-trained deep learning models stored on Zenodo (1), *WebIPred* allows clinicians to analyze patient data locally without transmitting it to external servers. In this study, we demonstrate *WebIPred*'s efficacy by using a pre-trained deep-learning model to detect pneumonia from chest X-ray images.

[1] Corresponding Author: Florian Auer; E-mail: florian.auer@informatik.uni-augsburg.de.

2. Methods

2.1. Dataset and Model Training

We utilized the Chest X-Ray Images (Pneumonia) dataset, which consists of anterior-posterior chest X-rays of pediatric patients aged one to five years from Guangzhou Women and Children's Medical Center (2). The dataset contains two classes: NORMAL and PNEUMONIA. All images were initially screened for quality control to remove low-quality or unreadable scans. Diagnoses were annotated by two expert physicians, and a third expert reviewed the evaluation set to account for any grading errors.

The dataset was split into training, validation, and test sets, and the model was trained using the AUCMEDI framework (3,4), a Python-based tool for medical image classification. AUCMEDI has been successfully applied to a variety of medical imaging datasets, achieving high classification performance (Table 1), and demonstrating strong accuracy in distinguishing between healthy and pneumonia-affected patients.

Table 1. Performance metrics of the *WebIPred* pneumonia classification model for PNEUMONIA and NORMAL classes.

Class name	Class frequency	Class percentage	Accuracy	AUC	F1
PNEUMONIA	1,583	27.0 %	98.5 %	99.9 %	99.0 %
NORMAL	4,273	73.0 %	98.5 %	99.9 %	97.4 %

2.2. Application Architecture

WebIPred is built using the Angular web framework (5) and relies on Tensorflow.js to load and apply pre-trained deep learning classifiers. The application runs entirely in the web browser, ensuring no patient data is transmitted to external servers. The web application is set up to download the trained model from Zenodo, a platform for sharing and preserving research data, including deep learning models.

Tensorflow.js requires the deep learning models to be stored in TensorFlow.js Layers format. This specific format consists of a JSON file that defines the model's architecture and a set of binary files containing the model's weights. Through Zenodo the models are provided in a compressed format (zip) to reduce the transmission load. This requires decompression within the web browser using the Javascript library JSZip. Once loaded, models are applied to medical images selected by the user, with all processing handled client-side through dedicated process workers. Angular web workers are a mechanism in Angular that allows for offloading CPU-intensive tasks to a separate thread, improving application performance and maintaining responsiveness.

2.3. Privacy Considerations

The primary design goal of *WebIPred* is to conform to stringent data protection regulations, such as GDPR in Europe, which impose strict limitations on how patient data can be handled. To achieve this, all image data is processed locally in the browser. *WebIPred* does not store any data on servers, and once a session is closed, all data - images, predictions, and model metadata - is deleted from the browser's memory. This architecture ensures that sensitive medical data remains secure, and there is no risk of data leakage to third-party servers.

2.4. Model Management and Loading

WebIPred supports a wide range of pre-trained models available from Zenodo or other repositories. These models must be compatible with the Tensorflow.js format. Users can filter models based on specific needs, and once selected, the model's binary files are loaded into the browser's runtime memory. The loaded model can then be applied to patient images, generating predictions directly on the client's device.

2.5. Prediction Workflow

WebIPred is designed to be user-friendly, even for clinicians with limited technical expertise. The prediction process is guided through a three-step form.

- **Model Selection:** Users select a classifier from the available models, with descriptions provided to facilitate decision-making.
- **Image Selection:** Users upload one or more medical images, which are validated for format compatibility.
- **Prediction and Visualization:** After confirming the selected model and images, *WebIPred* processes the images, generating class probabilities for each image. Results are visualized using bar charts and tabulated data and can be exported as CSV files for further analysis.

Figure 1. Prediction and visualization for image "person1345bacteria3425" from the Chest X-Ray Images data set within WebIPred. The image is with high probability correctly assigned to the pneumonia class.

3. Results

3.1. Web application

The application seamlessly handles image upload and preprocessing, ensuring compatibility with various image formats and automatically applying necessary transformations. *WebIPred* delivers accurate and interpretable predictions, leveraging efficient inference techniques and providing clear visualizations of results (Fig. 1). The application's ability to handle multiple images simultaneously and export results as CSV files enhances its usability for clinical workflows.

3.2. Privacy and Security

WebIPred's architecture ensured that no patient data left the local machine, making it fully compliant with data privacy regulations such as GDPR. This was confirmed through tests where all images were processed within the browser, and no data was transmitted to external servers. Once the user closed the session, all cached models and images were deleted, ensuring no residual data persisted.

3.3. Performance

WebIPred's reliance on client-side processing ensures privacy but introduces variability in performance depending on the user's hardware. In our tests, prediction times ranged from 2 to 20 seconds per image, with significant delays observed when using different web browsers and devices. Moreover, the current architecture does not cache models between predictions, requiring them to be reloaded for each session. This further impacts performance, especially when using large models. While *WebIPred* does allow for simultaneous processing of multiple images, the lack of model pre-heating results in noticeable delays, particularly when handling large datasets.

4. Discussion

WebIPred successfully addresses several key challenges associated with deploying machine learning in clinical environments. Its privacy-first design ensures that patient data remains secure, making it compliant with even the most stringent regulations. By running entirely in the client's browser, *WebIPred* avoids the need for any backend processing, minimizing the risk of data breaches and simplifying integration with clinical IT infrastructure.

However, interoperability in clinical systems is crucial for ensuring that different platforms and healthcare providers can seamlessly share and analyze diagnostic data. Generalizability to various deep learning models is equally important, allowing to adapt to different architectures and datasets, improving its reliability across diverse medical applications. Expanding support to medical image formats like Digital Imaging and Communications in Medicine (DICOM) and non-TensorFlow models would further broaden its applicability in clinical settings. This could be achieved by integrating the application with Picture Archiving and Communication Systems (PACS) to directly import DICOM images. Furthermore, it is crucial to encode clinician-validated diagnoses

into Electronic Medical Records (EMR) to ensure seamless integration into the clinical workflow.

Further research is required to optimize *WebIPred*'s performance. The client-side architecture introduces performance issues that may limit the application's scalability in busy clinical environments. The performance of *WebIPred* is heavily dependent on the client's hardware, and processing times can vary considerably. Solutions such as model caching, parallelization, and pre-heating could significantly reduce processing times and improve usability. However, this trade-off is still reasonable, as the added privacy protection gained from local inference outweighs the reduction in computational efficiency.

5. Conclusions

WebIPred represents a significant step toward the integration of machine learning into clinical workflows by providing a privacy-preserving, web-based application for medical image classification. In this study, we have demonstrated that it is possible to use high-performing models within the web browser for medical image classification tasks, exemplary for pneumonia detection from chest X-ray images. By leveraging pre-trained models from repositories like Zenodo, *WebIPred* allows clinicians to apply cutting-edge AI tools directly within their web browsers without compromising data privacy. While device performance can vary, most modern devices, can execute predictions within a reasonable timeframe.

An important feature of this framework is its privacy protection. Since the model runs entirely within the user's browser, no sensitive medical images need to leave the local machine, eliminating concerns about data transmission to external servers. This aspect is especially critical in healthcare applications, where data privacy is crucial. The potential to leverage powerful models in a privacy-preserving, client-side environment opens up new possibilities for secure, accessible, and scalable AI-driven healthcare solutions.

A live version with full functionality is hosted on GitHub Pages at https://frankkramer-lab.github.io/WebIPred and the corresponding source code for deploying own instances is available within the repository at https://github.com/frankkramer-lab/WebIPred.

References

[1] European Organization For Nuclear Research, OpenAIRE. Zenodo [Internet]. CERN; 2013. Available from: https://www.zenodo.org/

[2] Kermany DS, Goldbaum M, Cai W, Valentim CCS, Liang H, Baxter SL, et al. Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning. Cell. 2018 Feb 22;172(5):1122-1131.e9.

[3] Mayer S, Müller D, Kramer F. Standardized Medical Image Classification across Medical Disciplines [Internet]. 2022. Available from: https://arxiv.org/abs/2210.11091

[4] Schneider P, Müller D, Kramer F. Classification of Viral Pneumonia X-ray Images with the Aucmedi Framework [Internet]. 2021. Available from: https://arxiv.org/abs/2110.01017

[5] Jain N, Mangal P, Mehta D. AngularJS: A modern MVC framework in JavaScript. J Glob Res Comput Sci. 2014;5(12):17–23.